RESPONSE

This Response is filed in response to the Office Action mailed on March 4, 2005. At present, claims 1-4, 7-15, and 19-26 are pending in the Application. Claims 1-4, 7-15, and 19-26 stand rejected.

Examiner Interview

The Applicants thank the Examiner for speaking with them and their representative, Jason P. Fiorillo (Reg. No. 52,892) on June 23, 2005. Applicants also thank the Examiner for acknowledging that the pending claims are allowable over the cited art. As agreed during that telephone call, Applicants are not amending the claims in light of the cited reference, but are instead highlighting the distinguishing characteristics which make the pending claims patentable over the cited reference.

Rejections Under 35 U.S.C. §102

Claims 1-4, 7-15, and 19-26 were rejected under 35 U.S.C. §102(e) as being anticipated by U.S. Patent No. 6,772,368 to Dhong et al. ("Dhong"). The Applicants respectfully traverse the rejections as applied to the claims, and submit that Dhong does not anticipate, teach or suggest the claimed invention.

Dhong

Dhong is directed to a multiprocessor computer system with a "pair-wise high-reliability mode." (Title). The disclosed system is embodied by Fig. 3, in which "a high reliability processing system 300 is shown having first and second processors 305A and 305B." (Column 4, lines 29-30). As illustrated in Fig. 3, Dhong's system apparently comprises a single board on which processor 305A and processor 305B communicate through the use of a synchronous, on-board, bus. Notably, the compare unit 345 is also apparently on the same board. The operation of this system is set forth below:

> The system 300 includes a compare unit 345 for indicating whether the two processors' respective versions of instruction results match. If the versions do not match for an instruction, the instruction is deemed to be a faulting instruction.

> Responsive to the system 300 being in the high reliability mode and the compare unit 345 indicating a faulting instruction, the processors 305 recover a state that the processors had prior to execution of the faulting instruction, and the processors re-execute the faulting instruction. (Column 7, lines 8-16).

As mentioned previously, the two processors are apparently connected via a local, synchronous bus. This necessity is illustrated in Figure 3, and is described within the specification as follows:

> In the high reliability mode, the two processors 305 both execute the same program or instruction stream in parallel. That is, for one of the instructions in the stream each processor computes its own version of a result. In this mode, each processor 305 will need identical, lockstep bus 380 accesses, so mode control logic 302 notifies arbitration logic 348 in the bus interface unit 340 to allow only one of the bus logic units 344 to control bus 380 requests and read the bus 380 for both processors 305 in the system 300. Also, the arbitration logic compares requests from the processors 305, to check that both are generating the same requests in the same cycle. (Column 6, line 62 – column 7, line 7).

And further:

> Aside from synchrony in connection with the processors beginning execution of an instruction stream, another synchronizing issue also arises upon a demand for the processors to respond to an interrupt request. That is, the processors 305 may, of course, be subject to an external interrupt which could disturb synchrony, unless the dual processor 305 operation is coordinated properly. Accordingly, the processor 305 bus interface unit 340 has common external interrupt logic 342 which responds to external interrupt requests and signals logic 310 both processors simultaneously to respond to the interrupt request. The response may include merely setting a bit in a register for later follow up, or it may include causing the processor to branch to a micro code routine, execute a trap instruction calling an operating system routine, or even terminate dual execution of an instruction stream, so that the processors terminate in synchrony. That is, external signals to the BIU 340 that are going to the processors 305 are asynchronous. Since it would be impossible to guarantee that two arbitration circuits will always make identical decisions in synchronizing asynchronous signals, therefore this indeterminism is removed, according to an embodiment, by external interrupt logic 342 so that every external asynchronous signal is synchronized once, before it is copied to the processors 305, and no further synchronization is required in the processors 305. (Column 7, line 50 – column 8, line 8).

In sum, Dhong discloses a multi-processor system where two processors are connected on a single board by a local, synchronous, bus.

In contrast, Applicants' independent claims recite at least three novel elements not found in the Dhong reference. In particular, independent claim 1 recites "a fault-tolerant data processing apparatus comprising: a plurality of data processing elements executing substantially identical instruction streams substantially simultaneously; an I/O node in connection with at least one of the plurality of data processing elements; and a switching fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node." (Emphasis added).

A Plurality of Data Processing Elements.

As set forth in the specification and illustrated in Fig. 3, each data processing element comprises a "CPU board 22." Furthermore, "each CPU board 22 contains at least one processor 44 and the main memory 24. In some embodiments, each CPU board 22 contains multiple processors 44, 44', 44", and 44'" (generally 44). In multi-processor embodiments, each of the multiple processors 44 of a CPU board 22 may process different instruction streams. Respective processors 44 on different CPU boards 22 execute substantially identical instruction streams, and each processor 44 on a single CPU board 22 may execute its own instruction stream." (Paragraph 28). In addition, each CPU Board 22 may contain "a CPU bus interface 66, and CPU fault-tolerant logic 50. The CPU fault-tolerant logic 50 is in communication with the CPU bus interface 66. The priority register 46 of the I/O fault-tolerant logic 52 is in communication with the priority module 86 through a CPU PRIORITY signal 88. In some embodiments, the CPU fault-tolerant logic 50 on each redundant CPU 22 is in communication with the I/O fault-tolerant logic 52 through a respective command line 96, 96'." (Paragraph 41).

Thus, Dhong's "processors 305" are similar to the Applicants' "processors 44" but are distinguishable from the Applicants' "CPU Boards 22", or "data processing elements" as specified in claim 1. Accordingly, Dhong does not teach or suggest "a plurality of data processing elements executing substantially identical instruction streams substantially simultaneously."

A Switching Fabric.

Dhong discloses a local bus connecting two processors.

Conversely, Applicants disclose a "switching fabric" connecting a plurality of data processing elements. Applicants' switching fabric is illustrated in Fig. 6 of the application, and described as follows:

> The switching fabric 150 provides a means for communicating instructions and data between end nodes of the computer system 20. End nodes of the system 20 include CPU 22 end nodes and I/O end nodes 170, each associated with one or more peripheral devices 42.
>
> In one embodiment of a fault-tolerant computer system 20, the system includes a first and second CPU nodes 151 and 151' (generally 151), one or more I/O nodes 170a, . . . 170n, 170m (generally 170), the switching fabric 150, and a first and second voting module 153, 153' (generally 153). Each of the end nodes 151, 170 is in communication with the switching fabric 150 through at least one communications link 160. The switching fabric 150 includes network components, such as switches, routers, repeaters, and transceivers interconnected through communications links 160. The communications links 160 may be serial communication links, parallel links, or broadband links. The communications links 160 may be implemented with "twisted-pair" cable, coaxial cable, fiberoptic cable, or wireless links, such as radio links or free-optics links.
>
> In one embodiment the system includes at least one target channel adapter (TCA) 156 and at least one host channel adapter (HCA) 152. A TCA 156 is a channel adapter specifically configured for use at an end node including a peripheral device 42. An HCA 152 is a channel adapter specially configured for use at an end node including a CPU 22. In one embodiment, the switching fabric 150 is an INFINIBAND.TM. switching fabric 150 and the HCAs 152 and TCAs 156 are defined by the INFINIBAND.TM. standard. In another embodiment, the switching fabric 150 is an Ethernet switched fabric in a network using a transmission control protocol over internet protocol (TCP/IP) or other transmission protocol. In another embodiment, the switching fabric 150 is a packet switched fabric in a network using X.25, frame relay, ATM or other transmission protocol. In another embodiment, the switching fabric 150 is a circuit switched fabric, such as a wired circuit switched fabric in a wired network, a wireless circuit switched fabric in a wireless network, or a combination circuit switched fabric having elements of both wired and wireless circuit switched fabrics in a combination network. In yet another embodiment, the switching fabric 150 is an American National Standards Institute (ANSI) Fibre Channel switching fabric in a network using a fibre channel transmission protocol. (Paragraphs 55-57).

Thus, the Applicants' switching fabric refers to a connection which may be maintained over lengths longer than an on-board bus. In fact, Applicants' switching fabric may connect

computer systems located in geographically remote locations, as described above, via the use of Ethernet or across similar networks.

Accordingly, Dhong does not teach or suggest the use of a switching fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node.

Asynchronously.

As discussed previously, Dhong's on-board bus appears to operate in a synchronous, or lock-step fashion. In contrast, the Applicants' claimed invention makes use of a switching fabric which communicates transactions asynchronously between at least one of the plurality of data processing elements and the I/O node.

This asynchronous communication is described in the specification as set forth below:

> The transport messages, or packets, are communicated asynchronously across the switching fabric 150. Here, asynchronous means that the clock source of each of the end nodes 151, 170 operate independently from the clock sources of the other end nodes 151, 170, and also independently of the switching fabric clock source. Asynchronously also indicates that in some embodiments, data packets can be communicated across the switching fabric 150 intermittently where the intervals of any time delay between the communicated packets of the I/O transaction is variable. The variable delay and the nature of message routing where different packets of the same instruction may be routed through different paths within the switching fabric 150, thereby arriving at the TCA 156 in a different sequence than their sequence at the HCA 152. (Paragraph 65).

As mentioned above, Dhong does not teach or suggest the use of a switching fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node.

SUMMARY

Claims 1-4, 7-15, and 19-26 are pending in the application. Claims 1-4, 7-15, and 19-26 stand rejected. The Applicants request that the Examiner reconsider the application and claims in light of the foregoing Response, and respectfully submit that the pending claims are in condition for allowance.

If, in the Examiner's opinion, a telephonic interview would expedite the favorable prosecution of the present application, the undersigned attorney would welcome the opportunity to discuss any outstanding issues, and to work with the Examiner toward placing the application in condition for allowance.

Respectfully submitted,

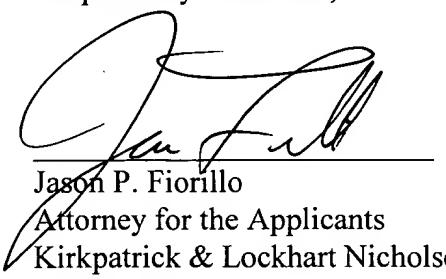

Date: June 24, 2005
Reg. No.: 52,892
Tel. No. (617) 261-3186
Fax No. (617) 261-3175

Jason P. Fiorillo
Attorney for the Applicants
Kirkpatrick & Lockhart Nicholson Graham
75 State Street
Boston, MA 02109

BOS-874397 v1